

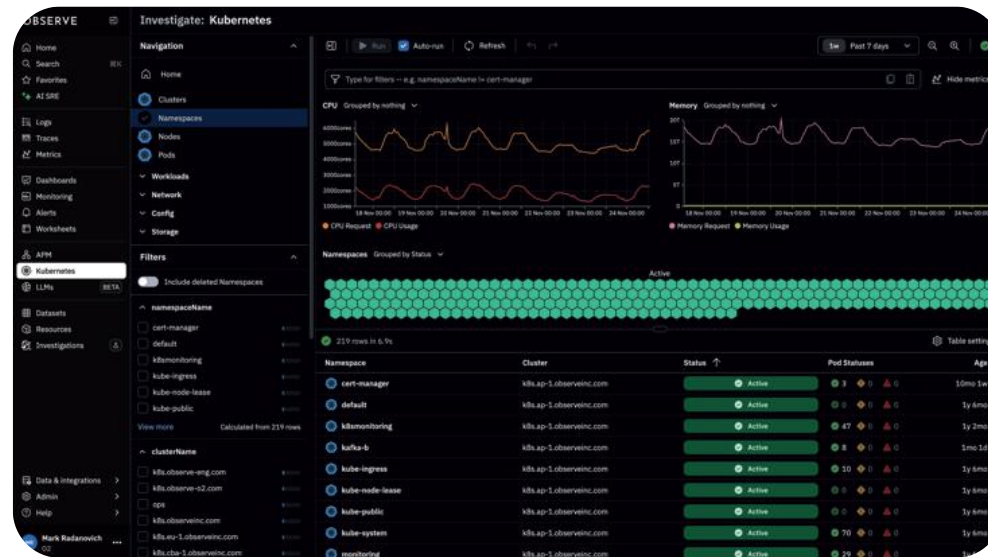

# Observe: Observability Engineered for Scale



AI-powered observability for 10x faster troubleshooting at 60% lower cost

## Scalable Observability for Logs, Metrics, and Traces

Modern systems generate more telemetry than legacy observability platforms can cost-effectively ingest. Observe solves the challenges of scale through a data lake architecture, which stores telemetry at dramatically lower cost, and a knowledge graph that tracks relationships across services and infrastructure. Pairing agentic workflows with context from the knowledge graph, Observe's chat-based AI SRE accelerates root cause analysis and resolution.



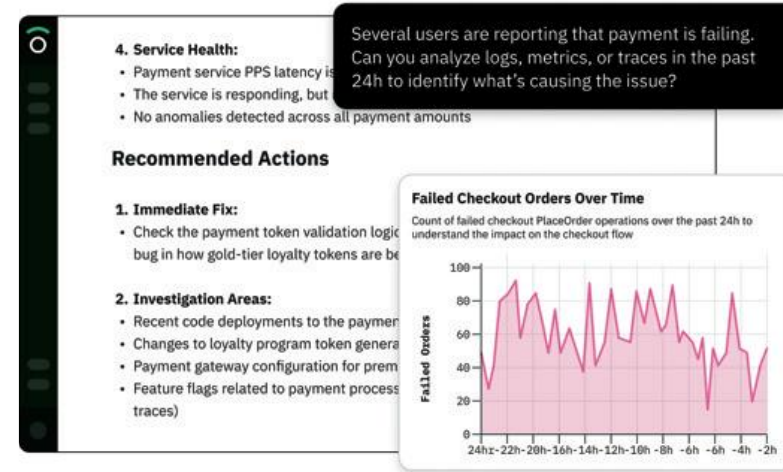
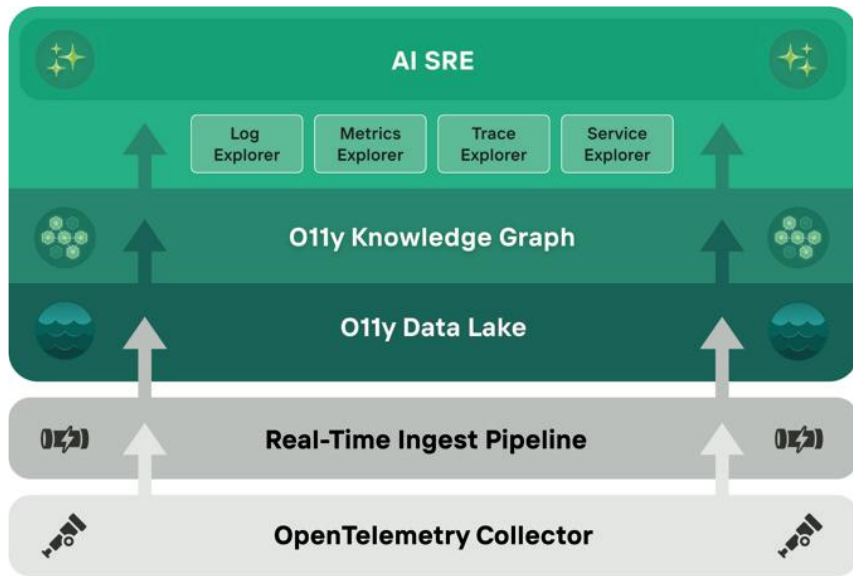
## TRUSTED BY



# Modern Observability Architecture



Observe’s cloud-native architecture maintains context across logs, metrics, and traces with a unified data lake and knowledge graph. Built on open standards like OpenTelemetry and Apache Iceberg, it ensures interoperability and avoids lock-in.



## Observe AI SRE

The AI SRE applies context, identifies root causes, and suggests fixes, enabling teams to reduce toil and troubleshoot 10x faster.

- Pinpoint failures and fix issues faster using contextual understanding of logs, metrics and traces
- Cut observability costs with a single low-cost data lake backend
- Operate at enterprise scale with built-in data privacy, compliance, and customizable agentic workflows via MCP Server

### AI SRE

Automates investigative heavy lifting to accelerate incident response and deliver precise root cause analysis.

### O11y Knowledge Graph

Structures and correlates telemetry for fast context-aware AI and analytics.

### O11y Data Lake

Keeps telemetry searchable longer at 60% lower cost. Supports Apache Iceberg for data sharing.

### Explorers

Interfaces for common data analysis use cases. Includes Log, Metrics, trace, Service and K8s Explorers.

### Elastic Compute

Scales compute as needed, without runtime interference from other workloads.